Alexandre Bouchard-Cote

Department of Statistics
University of British Columbia

Scalable approximation of integrals using non-reversible methods

December 4, 2019
11:45am – 12:45pm
Light refreshments will be served
110 Frelinghuysen Road
Hill Center, Room 552

Abstract: I will describe novel perspectives on the problem of approximating integrals, coming from the nascent field of non-reversible Monte Carlo methods. In particular, I will present an adaptive, non-reversible Parallel Tempering (PT) allowing MCMC exploration of challenging problems such as single cell phylogenetic trees.

By analyzing the behaviour of PT algorithms using a novel asymptotic regime, a sharp divide emerges in the behaviour and performance of reversible versus non-reversible PT schemes: the performance of the former eventually collapses as the number of parallel cores used increases whereas non-reversible benefits from arbitrarily many available parallel cores. These theoretical results are exploited to develop an adaptive scheme approximating the optimal annealing schedule.

My group is also interested in making these advanced non-reversible Monte Carlo methods easily available to data scientists. To do so, we have designed a Bayesian modelling language to perform inference over arbitrary data types using non-reversible, highly parallel algorithms.

Bio: My main field of research is in computational statistics/statistical machine learning. I am interested in the mathematical side of the subject as well as in applications in linguistics and biology. On the methodology side, I am interested in Monte Carlo methods such as SMC and MCMC, graphical models, non-parametric Bayesian statistics, randomized algorithms, and variational inference. My favorite applications, both in linguistics and biology, are related to phylogenetics in one way or another. Some examples of things I have currently/recently been working on: automated reconstruction of proto-languages; cancer phylogenetics; population genetics; pedigrees, tree and alignment inference. In the past, I also did some work on machine translation, on logical characterization and approximation of labeled Markov processes, and on reinforcement learning.